

Quiver: Using Control Perturbations to Increase the Observability of Sensor Data in Smart Buildings

Jason Koh, Bharathan Balaji, Vahideh Akhlaghi, Rajesh Gupta

University of California, San Diego
{jbkohl, bbalaji, vakhlagh, gupta}@eng.ucsd.edu

Yuvraj Agarwal

Carnegie Mellon University
yuvraj.agarwal@cs.cmu.edu

Abstract—Modern buildings consist of hundreds of sensors and actuators for monitoring and operation of systems such as HVAC, light and security. To enable portable applications in next generation smart buildings, we need models and standardized ontologies that represent these sensors across diverse types of buildings. Recent research has shown that extracting information such as sensor type with available metadata and timeseries data analysis is difficult due to heterogeneity of systems and lack of support for interoperability. We propose perturbations in the control system as a mechanism to increase the observability of building systems to extract contextual information and develop standardized models. We design Quiver, an experimental framework for actuation of building HVAC system that enables us to perturb the control system safely. Using Quiver, we demonstrate three applications using empirical experiments on a real commercial building – co-location of data points, identification of point type and mapping of dependency between actuators. Our results show that we can co-locate data points in HVAC terminal units with 98.4 % accuracy and 63 % coverage. We can identify point types of the terminal units with 85.3 % accuracy. Finally, we map the dependency links between actuators with an accuracy of 73.5 %, with 8.1 % and 18.4 % false positives and false negatives respectively.

I. INTRODUCTION

Cyber-Physical Systems (CPS) represent continuing advance of instrumentation of existing systems (such as automobile [1]) and networks (such as transportation [2], energy grid [3]) with a growing list of sensors and actuators. These sensing mechanisms (from tire-pressure meters to synchrophasers) provide a greater awareness of complete systems with increasingly finer resolutions of time and distance scales. Yet, deployment and maintenance of growing sensing infrastructure presents a significant challenge [4], [5], [6], one that researchers have tried to address through analysis of the sensory data. Actuation also provides us with increased observational capabilities by actively modulating a system and observing responses. Active control has shown to be effective across disparate disciplines – seismic structural design [7], aerodynamic stability analysis [8] and fault tolerant control [9]. However, control is seldom used across most systems and much of the literature is based on simulation studies. We propose that control perturbations be used across a variety of CPS applications such as information discovery, modeling, control optimization and privacy protection.

In this paper, we explore use of perturbation control to address sensor data management problem. Our thesis is that a carefully selected and controlled modulation of control

algorithms can be used to discover sensor conditions and reduce sensor maintenance work required for creation of CPS applications. We focus on buildings as a driver application to test this hypothesis. Buildings consist of hundreds to thousands of sensors and actuators that are used for management of air conditioning, lighting, fire safety and water. Prior work has addressed many challenges to enable development of smart building applications – standardized API for information access [10], data management [11], [12], and semantic ontology [13]. These solutions have led to innovative applications such as occupancy based control [14], [15], human in the loop control [16], [17] and energy disaggregation [18]. A major challenge in adoption of these applications is that they are not portable across different building systems due to lack of standardized ontologies [19], [20].

Recent works have focused on standardizing building information using existing sensor metadata and available timeseries data to enable portable applications [19], [20], [21], [22]. These works illustrate that modern buildings consist of a wide variety of sensors and actuators with varying naming conventions across vendors. Oftentimes the metadata available is not understandable to a non-expert and sometimes no metadata is available to understand the data context. Even when researchers used the available metadata, timeseries data and applied state-of-the-art algorithms to identify the *type* of the sensor, the overall accuracy was low [21], [22] or required significant manual effort [19], [20]. Identifying sensor type, however, is just one step towards creation of standardized models that can be used by portable applications. Other pertinent problems include determining sensor location, relationship among sensors/actuators and models which capture the behavior of the system.

Active control is a promising approach to address the lack of information available, as carefully designed control perturbations can reveal insights into system behavior that is not observed in regular operation. Recently, Pritoni et al. [23] showed that the mapping between the Air Handler Units (AHU) and the corresponding terminal units in the building Heating, Ventilation and Air Conditioning (HVAC) system can be inferred with 79% accuracy with control perturbations compared to 32% accuracy with data analysis alone. Control perturbations have also been studied for Fault Detection and Diagnosis (FDD) [24], [25] and fault tolerant control [26], [27] in HVAC systems as it eliminates mundane manual testing and

fixes some classes of faults automatically.

We expand on these ideas and show that active control mechanisms can be used as an integral part of a data model. Control based interventions are not used in practice because of equipment and safety issues. We empirically explore control in a real building HVAC system. We build Quiver, a control framework that allows us to do control experiments *safely* on the HVAC system by constraining control input that satisfies criteria such as range of values, frequency of actuation and dependency between actuators. We deploy Quiver in our building testbed and use it to demonstrate three example applications that exploit control perturbations. First, we show that perturbations can be used to identify co-located sensors which has been shown to be difficult with data alone [28]. We co-locate data points in HVAC terminal units with 98.4% accuracy and 63% coverage. Second, we identify the point type of terminal units given ground truth point types of one terminal unit using transfer learning classification. We identify point types with 85.3% accuracy across 8 zones. Third, we map the dependency between sensors and actuators in the control system using control perturbations and probabilistic analysis. We identify dependency links between actuators with 73.5 % accuracy, with 8.1% false positives and 18.4% false negatives across 5 zones.

II. OUR BUILDING TESTBED

Modern buildings consist of hundreds of networked sensors and actuators for operation and maintenance of various systems. These systems are typically overseen with Building Management System (BMS) which helps configure, monitor, analyze and maintain various systems. The sensors, actuators and the configuration parameters in the BMS are together referred to as *points*. We focus on building HVAC systems where BMSes are most commonly used.

Our testbed is a 150,000 sq ft building, constructed in 2004 and consists of a few thousand occupants and 466 rooms. The HVAC system consist of an Air Handler Unit (AHU) that supplies cool air to the building via ductwork using chilled water supplied by a central plant. A heat exchanger supplies hot water to the rest of the building using hot water supplied by central plant. The cool air and hot water are used by local terminal units called Variable Air Volume (VAV) boxes to regulate the temperature of rooms. The area serviced by the VAV box is referred to as a *thermal zone*, which consist of a large room or multiple small rooms in our building. Figure 1 shows a schematic of the VAV box with the sensors and actuators installed for its operation.

VAVs have been commonplace since 1990s [30], and their basic working is well understood. The VAV regulates the amount of cool air provided using a damper, and if the zone needs to be heated, it regulates the hot water in the heating coil using a valve. The temperature sensor in the thermal zone provides feedback on how much cooling or heating is required. However, in the real VAV box, there are over 200 points that govern its working [31]. The essential sensors include: *Zone Temperature*, *Supply Air Flow*, *Reheat Valve Position* and

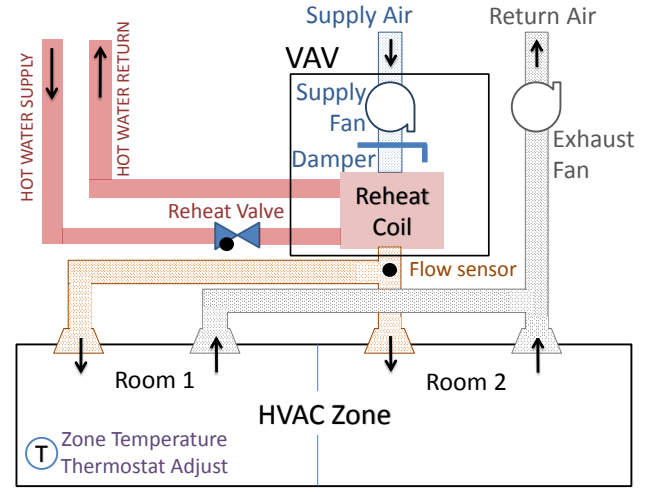


Fig. 1. Sensors and actuators in a Variable Air Volume (VAV) unit that provides local control of temperature in the HVAC system [29].

Damper Position; and the actuator points include: *Reheat Valve Command*, *Thermostat Slider Adjust* and *Damper Command*. These actuators are controlled using many configuration points such as *Temperature Setpoint*, *Occupied Command*, *Air Flow Setpoint*, etc. These configuration points account for majority of the points, and include nuanced parameters that ensure minimum airflow, set the PID loop settings, etc.

Not all of these 200 points are reported to the BMS, and only the essential sensors and control points are exposed to limit resource usage and information overload for building managers. In our building testbed, 14 to 17 points are reported to the BMS for each VAV box. The points exposed to BMS changes depending on the vendor, type of VAV and the installation version used by the vendor. Even though the same model of VAV is used across all zones in our building, there are minor variations due to configuration changes, presence of supply/exhaust fans or lack of heating.

A. Data Collection and Control

The points in our building communicate with the BMS using BACnet [32], a standard building network protocol. We connect our server to this network to collect data and control the points in our building. We use BuildingDepot [33], an open source RESTful web service based building datastore to manage the points in the building, provide appropriate permissions to developers, and search using a tagging mechanism. Our control framework Quiver works on top of BuildingDepot to manage control inputs from our experiments. Figure 2 depicts the system architecture of our deployment.

BACnet is a well developed protocol with which developers can not only read and write points, but also schedule hourly control, mark holidays on a calendar, and even manage programs running in the embedded VAV controller. For simplicity, we only focus on read and write points, i.e., in BACnet terminology *Input*, *Output* and *Value* points. These points can have floating point, binary or multi-state values, and in BACnet

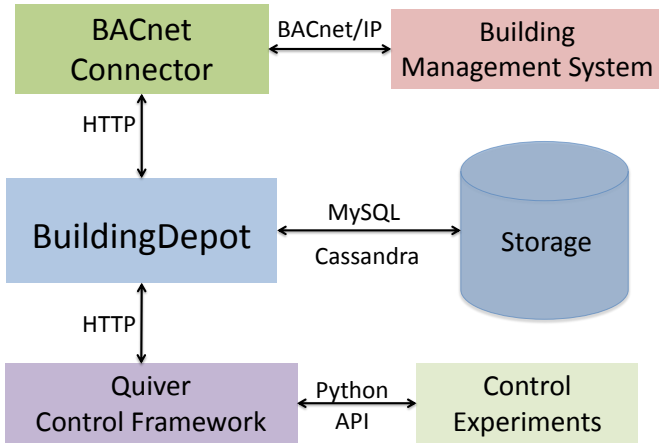


Fig. 2. System architecture of Quiver. Data collection and control is done via BACnet protocol using BuildingDepot web service [33]. Quiver ensures that the control sequences of our experiments are safe and rolls back the system to its original behavior in case of failure.

a floating point that can be written to is referred to *Analog Output*. Each of these *Output* points have an associated priority array. The default operation is performed at the lowest priority and the highest levels are reserved for emergency operations such as fire safety. Once a higher level priority is written to, the lower levels are ignored. An explicit write with value “0” needs to be written to the higher level priority in order to relinquish control to the lower levels.

The university Facilities Management provides us with a fixed priority level in this priority array for our control experiments. We need to relinquish control back to the default priority level after our control experiments to ensure that our interference does not affect the regular operation of the HVAC system. Quiver ensures that all the points are relinquished after an experiment.

B. Points in Variable Air Volume Box

Figure 3 shows the points associated with VAV in our building BMS and how these points relate to each other. At the top of the figure, we have the zone *Temperature Setpoint* and *Occupied Command*, which in combination with thermostat input determine the temperature guardband within which the VAV is trying to keep the zone temperature. The temperature guardband is indicated by *Heating* and *Cooling Setpoints*, which represent the lower and upper bounds of temperature respectively. There are three occupancy modes: *Occupied*, *Standby* and *Unoccupied* during which the temperature bands are $4^{\circ}F$, $8^{\circ}F$ and $12^{\circ}F$ respectively. During the *Occupied* mode, minimum amount of airflow is maintained to ensure indoor air quality. The *Thermostat Adjust* allows changing the temperature setting by $\pm 1^{\circ}F$, and the *Temporary Occupancy* maps to a button on the thermostat which when pressed puts the zone to *Occupied* mode for two hours during nights/weekends.

The *Heating* and *Cooling Setpoints* determine the behavior of the VAV control system with the measured *Zone Tem-*

perature completing the feedback loop. These three points determine the *Cooling* and *Heating Command* of the thermal zone. The *Cooling Command* determines the amount of cool air required for the zone and determines an appropriate *Supply Air Flow Setpoint* that is between the designed minimum and maximum supply air flow. When the *Cooling Command* is high ($\sim 100\%$), feedback is sent to the AHU to decrease the supply air temperature to meet the cooling needs of the thermal zone. The *Heating Command* determines the amount of reheat required by controlling the *Reheat Valve Command*. During heating, the airflow is set to the minimum to reduce chilled airflow from AHU, and this airflow is increased when high *Heating Command* ($\sim 100\%$) fails to heat up the thermal zone sufficiently. A high *Heating Command* also sends a signal to the heat exchanger to increase the supply water temperature. Note that only one of *Heating* or *Cooling Commands* can be $>0\%$ at a time.

The *Supply Air Flow Setpoint* determined by the cooling/heating requirements in turn determines the *Damper Command* which is the amount of damper actuation required to match the setpoint to the measured *Supply Air Flow*. The *Damper Position* sensor also provides feedback to set the appropriate *Damper Command*. There is a separate PID loop associated with setting each of *Heating Command*, *Cooling Command*, *Supply Air Flow Setpoint* and *Damper Command*, and there are PID parameters such as those that govern proportional gain and integration time, but these are hidden from the BMS.

III. QUIVER: CONTROL FRAMEWORK

Quiver¹ is built upon our building testbed. In this section, we show the utility of Quiver on the VAV box to demonstrate our control perturbation applications.

Given that exercising control over a building control system can lead to unintended, and potentially dangerous consequences, we ensure safety for the control perturbations we perform by: (a) global time synchronization across the computational servers running different parts of Quiver, (b) identification of a range of ‘safe’ values for control through metadata and data history, (c) dependency checks based on domain knowledge, and (d) status tracking to ensure we relinquish control at the end of our experiment and are able to rollback the changes made in case of an application crash. We have designed these safeguards based on five years of control experience with BMSes and discussions with our university Facilities Management.

We perform basic checks such as checking the range of values exercised (i.e. min/max), type of data for each input based on metadata available in BACnet as well as self-imposed limits. For instance, we force the *Temperature Setpoint* to be between $62^{\circ}F - 78^{\circ}F$. Each point marked as read/write in Figure 3 can be written to using Quiver. We synchronize time across all our servers – BACnet Connector, BuildingDepot and

¹The basic control framework will be presented as a poster abstract (not peer reviewed) in an upcoming conference.

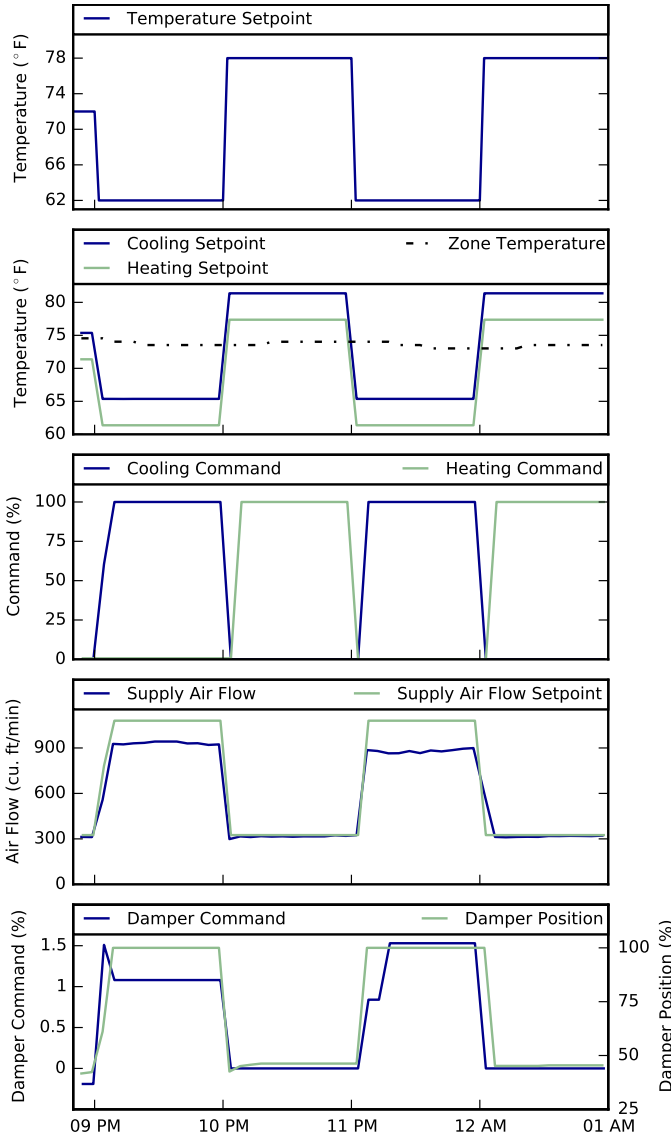


Fig. 4. A sample of co-location experiment, where the *Temperature Setpoint* is oscillating between $62^{\circ}F$ and $78^{\circ}F$ for four hours at night (the top graph). The VAV points which react to these changes (remaining four graphs) can be co-located by using temporal data analysis of this controlled period.

- Identifying points which are co-located with a VAV box.
- Identifying point types within a VAV given ground truth point types of one VAV unit.
- Mapping the dependency between VAV actuator points.

All of our data analysis is implemented using Python Scikit Learn library [34].

A. Using Control Perturbation to Determine Location

The location of sensor and actuator points is not readily available in the BMS for older buildings. In buildings where location information is available, it is usually inconsistent due to errors in manual labeling process [20], [21]. It is also difficult to co-locate points using historical data alone as many VAVs function in a similar manner, and the variation of data is not enough to distinguish them apart [28]. Control

perturbations can be used to force the control system to unusual operating points, and co-located points that respond to this perturbation can be clustered together by data analysis.

We assume that we already know the type of points in the building which can be obtained using recently proposed methods [20], [22], but do not know if they are co-located or how these point relate to each other or affect the control system. We do not use the location information already integrated into Quiver for these experiments. We perturb the actuator point identified as the *Temperature Setpoint* (TS) of a randomly chosen zone, and identify the corresponding co-located points using the temporal data features of other points. Towards the end of this section, we discuss how we can relax the assumption of knowing the point type apriori.

Figure 4 shows an example control sequence, where we change TS 4 times across 4 hours from low ($62^{\circ}F$) and high ($78^{\circ}F$), and the corresponding VAV points in the same zone that react to its changes. We chose such an oscillation of TS as it deviates substantially from normal operation and we can easily distinguish the controlled zone from the rest of the zones under normal operation. This control sequence was chosen empirically, and we show that even such simple control sequences can be effective for co-location of points. However, as we show with our experiments, the effect of control sequences do affect the quality of results. We do not focus on design of generic control sequences in this paper.

We extract basic features such as range (max - min), mean and standard deviation from the observed timeseries data. We also extract the Dynamic Time Warping (DTW) distance [35] between the applied TS signal and the point under consideration. DTW compensates for the time delay in the reaction and change in sensor values due to a control action and quantifies the difference between the shape of the signals. In addition, we exploit our pulse control and analyze the Fast Fourier Transform (FFT) after normalizing the data and use the Euclidean distance between the FFT of the point signal and FFT of the TS signal. We refer to this feature as “L2 norm of FFT” or “LFT”. We ignore frequencies beyond 0.0005 Hz, i.e. a period of 30 minutes, because we only focus on changes caused by our low frequency control sequence.

We extract these features for all the VAV points in the building and identify the outlier points. In principle, the point which deviates the most from regular control operations would be co-located with our TS point with high probability.

Figure 5 shows the distribution of all the *Zone Temperature* (ZT) points in our building across three features – DTW, LFT and range – with a control sequence of two changes to TS. The zone under control is marked in red, and as observed, the red point is far away from most of the points from the other zones in the building. However, there are still a few points which are also differ significantly from most zones and it is difficult to distinguish the red point from those outliers. When we examine the data from the experiment where we made 4 changes to TS, the corresponding changes to the ZT in the same zone resulted in much higher variation from those in uncontrolled zones. This is captured by our features as shown

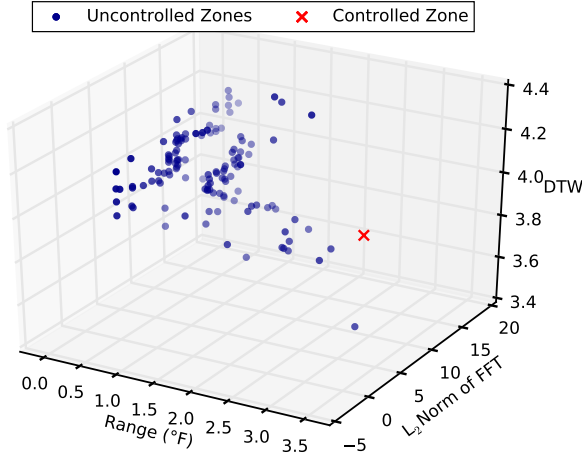


Fig. 5. Co-location of Zone Temperature by perturbing the Temperature Setpoint. 2 changes of Temperature Setpoint is applied across 2 hours in the controlled zone.

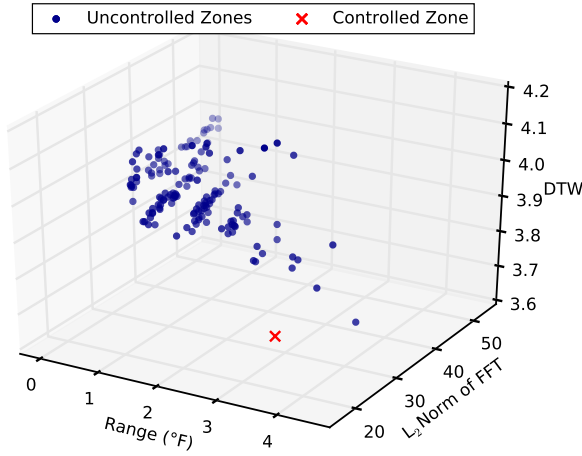


Fig. 6. Co-location of Zone Temperature by perturbing the Temperature Setpoint. 4 pulses of Temperature Setpoint is applied across 4 hours in the controlled zone.

in Figure 6. Hence, with the help of a well designed control perturbation it is possible to mold the behavior of the control system for end use applications.

We analyze the data for other point types to check if we can co-locate the zonal points successfully. In practice, we find that LFT feature alone is sufficient to distinguish the controlled zone points from the rest. Figure 7 shows compares the LFT of the controlled zone with other zones for point types: *Zone Temperature* (ZT), *Supply Air Flow Setpoint* (SAFS), *Supply Air Flow* (SAF), *Reheat Valve Command* (RVC), *Heating Command* (HC), *Damper Position* (DP), *Cooling Command* (CC), *Heating Setpoint* (HS) and *Cooling Setpoint* (CS). We performed this control experiment on eight zones in our building, and we co-located the listed points with 98.6% accuracy. We only failed to identify the correct *Damper Position* point

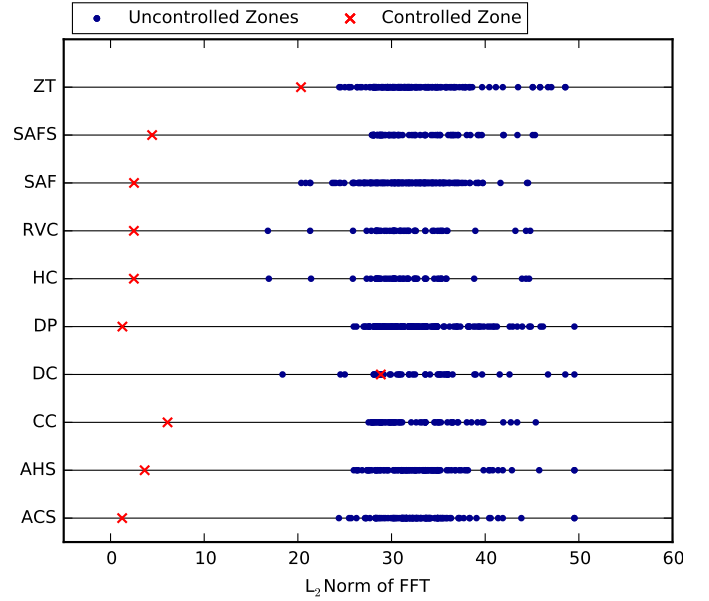


Fig. 7. Comparison of the L2 norm between FFT of VAV points and the FFT of the controlled *Temperature Setpoint*. The points corresponding to the controlled VAV have a much lower L2 norm compared to regular zones for eight point types, but fails to capture the difference in *Damper Command*.

for one of the zones, leading to a drop in accuracy.

The *Damper Command* (DC) is a differential actuator that sets the change that needs to be made to the damper. There are several VAVs in the building which constantly change their DC for minor variation in the air flow, and the features we extracted – DTW, FFT, mean, variance, number of changes – failed to differentiate the DC of the zone under control from the rest (Figure 7). More sophisticated data analysis or perturbation signals are required for co-location of DC points. We could only co-locate DC points in two of the eight zones with our current method.

Another issue with these control experiments is that we can only co-locate those points which react to changes in TS (see Figure 3). Points such as *Occupied Command* (OC) and *Thermostat Adjust*, which are external inputs to the VAV control system, cannot be co-located. To remedy this, we perform a second set of experiments which oscillates the OC similar to our TS control perturbations. We successfully co-located OC and points such as *Cooling Setpoint* and *Supply Air Flow* that have been already co-located with their corresponding TS point. Thus, all of these points can be marked as being co-located in the same VAV. We performed the *Occupied Command* oscillation experiments across four zones with 100% success rate in their co-location results. We could not perform similar experiments on the thermostat points (*Thermostat Adjust* and *Temporary Occupancy*). They cannot be controlled by our platform as thermostats produce their data contiguously, and we acknowledge this is a limitation of our proposed method.

We repeated our TS control experiments on a hot day, and found that the same control perturbations cannot co-locate

heating related points – *Heating Command* (HC), *Reheat Valve Command* (RVC) – as they are not triggered sufficiently due to hot outside weather. The zone cannot be cooled down enough to activate heating (HC, RVC) when its TS is changed to the high value. We need to change our control perturbations to excite these points specifically. Thus, the perturbation signature needs to be sensitive to external conditions and confounding factors.

Note that in Figure 7, the LFT feature of the controlled zone for most point types except ZT and DC differ significantly from the rest of the points. Controlled zone’s LFT feature of ZT is not distinguishable from other types’ LFT feature of the other zones though it can be co-located within the same type. Contrary to DC’s different operational behavior, ZT’s signal response is slower than the other types due to heat capacity of zones. If we assume that we do not know the point type, then the points except ZT and DC can be obtained as outliers from the points belonging to normal zones. However, we would need to design appropriate threshold or clustering technique to identify the outlier points correctly.

Overall, with our control perturbations based co-location, we successfully co-located 10 out of 16 point types (63% coverage) with 98.4% accuracy across eight VAV units. We co-located the *Occupied Command* points using auxiliary control experiments for four zones. We failed to co-locate *Damper Command* due to its divergent behavior. We also do not co-locate 4 point types which do not respond to control perturbations.

B. Identification of Point Type

We now look at the inverse problem of identifying point type given the co-located points in a VAV. Identification of point type is essential for third party application to interpret timeseries data, and prior work has shown that data analysis alone fails to identify point types accurately [21], [22].

As shown towards end of Section IV-A, it is possible to identify co-located points even when the individual point types are unknown. For this problem, we assume that we know the co-location of all the VAV point types and the ground truth point types for one zone. We use this information along with the timeseries data to identify point types of other zones in the building. When our data analysis fails to identify certain point types, we use control experiments to increase the coverage of point types identification. Note that the VAV points we focus on cover 79% of all points in our building testbed.

With the help of the zone for which ground truth is known apriori, we train supervised classifiers using features extracted from one year of timeseries data and ground truth point types as labels. We slice one year of data into 53 weeks (partial weeks for starting and ending week), and extract features such as mean, variance, dominant frequency, noise, skewness and kurtosis [36] for each week. Noise is represented by error variance between original data and its piece-wise constant approximation. Skewness measures the symmetry of the data values across the mean while kurtosis measures the peak of the data distribution at the mean. The vector of

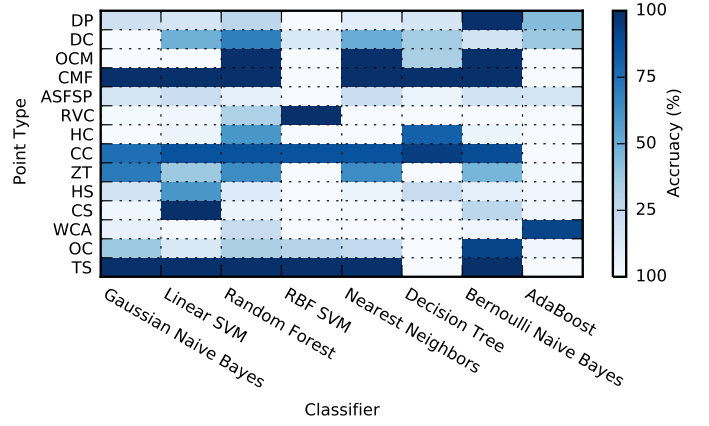


Fig. 8. Accuracy of point type identification using transfer learning across ten supervised classifiers with one year historical data for VAV points belonging to one zone.

features for one week of data represents one row of training data for our classifier. We train many standard classifiers – Gaussian Naive Bayes (GNB), Linear Support Vector Machine (LSVM), Random Forest (RF), Radial Basis Function SVM (RBF SVM), Nearest Neighbours (NN), Decision Tree (DT), Adaboost, and Bernoulli Naive Bayes (BNB). Our test results show high accuracy ($\sim 100\%$). We use these trained models for identifying point types of *different* zones with the same features extracted from one year of data. This process is called *transfer learning*, and figure 8 summarizes the results obtained across all the classifiers for different point types in one zone.

We find that some point types such as *Temperature Setpoint* (TS) are readily identified even with simple features such as mean and variance across most classifiers, while other point types work only with specific classifiers or require more features. We experimented with 10 classifiers for three zones, and choose the classifier which works best for each point types for analysis with other zones. Some point types such as *Heating Command* (HC), *Cooling Command* (CC) and *Zone Temperature* (ZT) are not identifiable by any classifier with all the features we used. This is because the behavior of these point types are similar to each other or other identified point types during regular operation. We repeated this analysis across 8 zones and found similar results. Our results are summarized in Table I. These results are corroborated by prior work who used more sophisticated data features for type identification [19], [21], [22].

We leverage control perturbations to identify point types which were unsuccessful using data analysis. As points such as TS and *Occupied Command* (OC) are easily identifiable using data analysis, we use these points to create our control perturbations. We put the zone to the *Occupied* mode, and increase the TS to $78^\circ F$ to force the VAV into a heating mode for 3 hours. The same perturbation is applied for the ground truth zone as well as all the zones for which the point type needs to be identified. We then extract the same data features for this control period across all the points, and use transfer learning to label points in different zones. As the behavior of

TABLE I
RESULTS OF POINT TYPE IDENTIFICATION USING BOTH DATA ANALYSIS
AND ANALYSIS WITH CONTROL PERTURBATIONS.

Point Type	Data		Control	
	Classifier	Accuracy	Classifier	Accuracy
Temperature Setpoint	RF	100%	-	-
Occupied Command	RBF SVM	99.2%	-	-
Thermostat Adjust	AdaBoost	93.1%	-	-
Cooling Setpoint	LSVM	88.9%	-	-
Heating Setpoint	DT	87.4%	-	-
Zone Temperature	RF	65.0%	RF	88%
Min Supply Flow	RF	77.8%	-	-
Max Supply Flow	BNB	100%	-	-
Cooling Command	DT	47.0%	DT	88%
Heating Command	DT	32.1%	-	-
Reheat Valve Cmd	RBF SVM	78.4%	-	-
Supply Air Flow SP	NN	80.3%	-	-
Supply Air Flow	LSVM	90.6%	-	-
Damper Command	RF	96.2%	-	-
Damper Position	RF	80.3%	-	-

all the controlled zones is forced to be similar, the behavior of the same point type across the zones is similar. However, the controlled sequence also forces the non-identified point types – SAF, SAFS, ZT, HC – to be different from the rest of the points. Thus, with this experiment, we are able to identify the rest of the point types successfully. Our results across the 8 zones have been summarized in Table I. Overall, our accuracy of point type of identification is 85.3 %.

Our data analysis or control experiments could not distinguish between point types *Heating Command* (HC) and *Reheat Valve Command* (RVC) as they are virtually identical in their behavior. However, it is possible to separate these two points using dependency analysis, which we describe next.

C. Dependency Mapping of Points

We now focus on the understanding the working of the VAV, and how the points relate to each other. As the type of points exposed are different across vendors and equipment, it is necessary to understand the context of these points, and map it to a model that can be used by other applications. These models can be built using domain knowledge, technical documents and historical data analysis [37], [38] as demonstrated by our dependency graph in Figure 3. We propose control perturbations as an alternative to these methods, which can be used for either verifying already developed models or used for older buildings where information available is insufficient for modeling using other methods.

We assume that we already know the point type and the co-located points in a VAV. We focus on modeling the dependencies between the actuator points (or read/write points).

We write to the *Temperature Setpoint* (TS) of a zone with a randomly chosen value every 20 minutes for 6 hours. The goal of this control sequence is to identify the points that react to changes in TS, and so we choose random values within limits

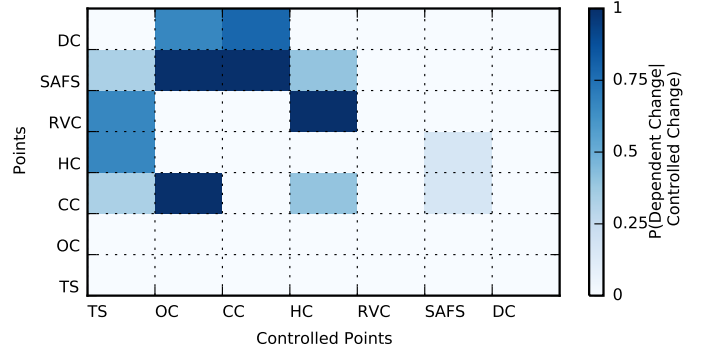


Fig. 9. Color map showing the changes induced by control perturbations of each actuator point on other actuators. Probabilities are calculated as the ratio of number of changes observed in a non-controlled actuator and the number of changes made by the controlled actuator.

for perturbing different operating points of the control system. For every change in TS, we analyze the behavior of the VAV points for 10 minutes, and note the points whose values change during this period. The threshold of change is one standard deviation for values observed for the past 12 minutes. We chose these times so that we can isolate changes that occur due to the change in TS rather than other external factors such as solar radiation. For the duration of the experiment, we calculate a final probability for each of the points as the ratio of the number of changes observed for the point and the number of changes in TS. We repeat this experiment by perturbing all the actuator points - *Occupied Command* (OC), *Cooling Command* (CC), *Heating Command* (HC), *Supply Air Flow Setpoint* (SAFS) and *Damper Command* (DC). Note that we ignore the points related to minimum and maximum supply air flow as they are constants.

Figure 9 shows a color map representing the probabilities obtained by perturbing each of these points. The changes to TS affects all the actuators except OC and DC, while changes to actuators like CC cause changes only in DC and SAFS. With the help of this color map, we can understand which points are being affected by each of the actuator points. However, this does not precisely decide the dependency between points as points which are lower in the dependency tree such as SAFS get affected by almost all of the actuation experiments. We find the behavior of DC to be unpredictable, and the changes that occurred in DC with our control perturbances were lesser than our set threshold. We perform these perturbation across five zones.

Figure 10 shows the relationships obtained as a result of our analysis. The green solid links indicate relationships which are true and confirmed with the analysis. The red dashed links show relationships which are not true, but are shown to be related by the analysis. In general, the above experiments cannot identify the true links when a “cycle” is formed in the graph. Here, by “cycle” we mean that there are multiple paths from one point to another in the graph. We perform more control experiments to negate the red links, and also confirm the blue dotted links which form a cycle with the green links.

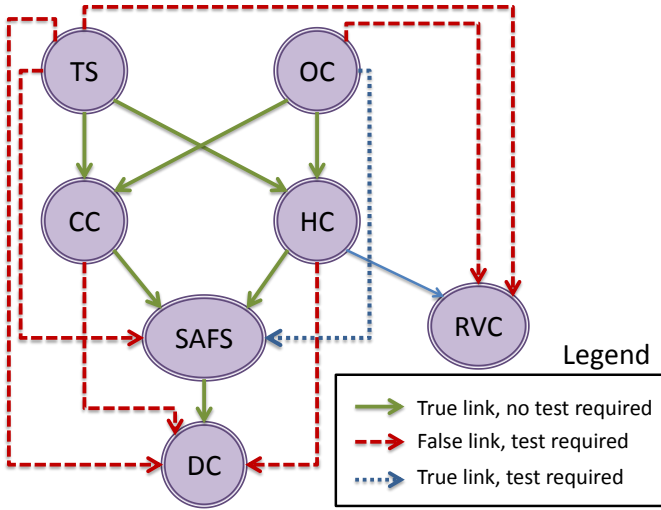


Fig. 10. Dependency links obtained by perturbing each of the actuators and checking which other actuators react to these changes. The links which require further testing are verified using conditional control perturbations.

Consider the cycle formed between the points TS, CC and SAFS. In order to verify if the link between TS and SAFS is correct, we perform a conditional control perturbation, where we change TS but force CC to be unchanged, which is called *Graph Surgery* [39]. We repeat this experiment for at least four changes of TS. If TS were directly affecting SAFS, we would observe that SAFS changes even when CC is held constant. We verify each of the red and blue links this way. When there are more than three points in a cycle, such as that with TS, OC, HC and RVC, we ensure several combinations of TS and OC are performed to test validity of the link. In some of these experiments, we preset the TS value to a fixed value for appropriate conditions that can activate other points such as HC. As we note at Section III, an external variable, ZT, may disable HC or CC in certain condition though ZT is not an actuator. We let the VAV control system to settle to a steady state after our change of TS before performing any dependency experiments.

We performed these experiments on five zones in our building, and verified the links with 73.5 % accuracy with a false positive rate of 18.4 % and false negative rate of 8.1 %. All of the false positive and 64.7 % of the false negative are due to the external variable, ZT. Thus, we can discover the dependency between actuator points in the VAV using our control experiments. However, in order to discover the complete dependency map as shown in Figure 3, we need to use maximum likelihood data analysis as the behavior of read-only points cannot be controlled directly.

V. RELATED WORK

The problem of discovering system characteristics and models using available data is called *system identification* [40], and is a well studied subject in control systems research. Using control perturbations for system identifications is also well known [40], [41], and the design of control perturbations,

also referred to as auxiliary or secondary signals, has been studied for modeling different types of control systems [42], [43]. System identification techniques have also been used for HVAC modeling [44], and some prior work have explored using control signals for system identification [45], [46]. However, all of these works focus on modeling the control system or perform control optimizations and do not address identification of contextual information such as location, point type or dependency graphs. Moreover, the control perturbation methods used for buildings are only verified using simulations.

Active HVAC control on real systems has been used for fault diagnosis [24] and fault tolerant control [26]. These works and other simulation based studies which propose fault tolerant control [47], [48] assume the contextual information about the system is already available. We focus on discovering contextual information using active control.

Co-location of sensors has been studied before [49], [28] but they use sophisticated data analysis algorithms. These methods fail when the points from different locations have similar data characteristics during regular usage. We show that with perturbations we can excite the local control system to unique operating points and co-locate points with high accuracy using simple data analytics.

Point type identification has also been studied earlier using both metadata [19], [20] and data analytics [21], [22]. These works show that metadata alone can be unreliable and requires significant manual input for accurate type identification, and the data analytics based method is useful for some point types but fails for others. Our results conform to the data analysis works, and we show that perturbations can be used to identify the points which are difficult with data analysis. We also note that Hong et al. [22] focus on transfer learning across buildings, while we focus on transfer learning within the building across the different instances of VAV units.

In addition, we focus on creating dependency graphs between points in the HVAC system. Pritoni et al. [23] use control of AHU discharge air temperature to map the VAV units to their corresponding AHUs. We present control techniques to isolate relationships within the VAV unit.

We have presented a detailed view of how HVAC VAV units work in practice, and how our Quiver control framework has been designed for safe control of the VAV unit, and provides features such as query of current status and rollbacks in case of crashes. Dawson-Haggerty et al. [50] provide similar mechanisms to provide safe operations, but the onus of rollbacks and error checking is on the application developer. They also provide support for multi-user control with timed leases and locks, while we only support single user control.

VI. LIMITATIONS AND FUTURE WORK

We have shown that control perturbations can be useful in discovery of contextual information in HVAC VAV units in our building testbed. This is just a first step, and more research is needed to extend these results to other types of equipment and control systems. We hope that the research community embraces control perturbations as a tool and generalizes the

results across different buildings, vendors and even other cyber physical systems.

Although our control perturbations were simple and effective for extracting information we were interested in, we observe that design of perturbation signatures can affect the results significantly depending on external disturbances and behavior of the control system. We will focus on automating the perturbation signals in future work, borrowing ideas from system identification literature [43]. We have also forced our control experiments to be done on either weekends or nights. Our work needs to be extended to perform these experiments even during work hours with perturbation signals that conform to comfort constraints.

To fully exploit control capability for both perturbation applications and control optimizations, we need to provide support to third party developers for direct control of systems. Initial work in both academia [50] and industry [51] have extended traditional vendor specific controls to common APIs. Our control framework, Quiver, allows safe control by a single user. Further work needs to be done to provide safety guarantees so that multiple users can access the control system at the same time and developers are provided sandboxes or emulation tools to experiment with different control applications.

VII. CONCLUSION

Traditional Building Management Systems are vendor specific vertically integrated solutions. To enable third party building applications that exploit sensor information, recent works have created standardized APIs and data management solutions which have led to a spurt of innovative applications. We extend these works by enabling control of building systems and demonstrate applications that exploit control to extract contextual information. We design Quiver, a control framework that allows safe experimentation of HVAC VAV units by integrating domain knowledge gained through experience and empirical experiments. Using Quiver, we demonstrate three control perturbation applications that extract context information about VAV units – co-location of points, identification of point types and mapping the dependency between points. We co-locate 63% VAV points with 98.4% accuracy, we identify point types with 85.3 % accuracy across 8 zones, and we identify dependencies between VAV actuator points with 73.5 % accuracy across 5 zones.

REFERENCES

- [1] C. Berger and B. Rumpel, "Autonomous driving-5 years after the urban challenge: The anticipatory vehicle as a cyber-physical system," *arXiv preprint arXiv:1409.0413*, 2014.
- [2] H. Abid, L. T. T. Phuong, J. Wang, S. Lee, and S. Qaisar, "V-cloud: vehicular cyber-physical systems and cloud computing," in *Proceedings of the 4th International Symposium on Applied Sciences in Biomedical and Communication Technologies*. ACM, 2011, p. 165.
- [3] S. Karnouskos, "Cyber-physical systems in the smartgrid," in *Industrial Informatics (INDIN), 2011 9th IEEE International Conference on*. IEEE, 2011, pp. 20–23.
- [4] P. Derler, E. Lee, A. S. Vincentelli *et al.*, "Modeling cyber-physical systems," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 13–28, 2012.
- [5] S. Karnouskos, "Stuxnet worm impact on industrial cyber-physical system security," in *IECON 2011-37th Annual Conference on IEEE Industrial Electronics Society*. IEEE, 2011, pp. 4490–4494.
- [6] B. Balaji, A. Faruque, M. Abdullah, N. Dutt, R. Gupta, and Y. Agarwal, "Models, abstractions, and architectures: the missing links in cyber-physical systems," in *Proceedings of the 52nd Annual Design Automation Conference*. ACM, 2015, p. 82.
- [7] M. D. Symans and M. C. Constantinou, "Semi-active control systems for seismic protection of structures: a state-of-the-art review," *Engineering structures*, vol. 21, no. 6, pp. 469–487, 1999.
- [8] A. Epstein, J. Ffowcs Williams, and E. Greitzer, "Active suppression of aerodynamic instabilities in turbomachines," *Journal of Propulsion and Power*, vol. 5, no. 2, pp. 204–211, 1989.
- [9] M. Blanke and J. Schröder, *Diagnosis and fault-tolerant control*. Springer, 2006, vol. 2.
- [10] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler, "sMAP: A simple measurement and actuation profile for physical information," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2010, pp. 197–210.
- [11] P. Arjunan, N. Batra, H. Choi, A. Singh, P. Singh, and M. B. Srivastava, "SensorAct: A privacy and security aware federated middleware for building management," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 80–87.
- [12] A. Krioukov, G. Fierro, N. Kitaev, and D. Culler, "Building application stack (bas)," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 72–79.
- [13] Project Haystack. <http://project-haystack.org/>.
- [14] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, "Sentinel: occupancy based hvac actuation using existing wifi infrastructure within commercial buildings," in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2013, p. 17.
- [15] P. X. Gao and S. Keshav, "Spot: a smart personalized office thermal control system," in *Proceedings of the fourth international conference on Future energy systems*. ACM, 2013, pp. 237–246.
- [16] V. L. Erickson and A. E. Cerpa, "Thermovote: participatory sensing for efficient building hvac conditioning," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 9–16.
- [17] A. Krioukov, S. Dawson-Haggerty, L. Lee, O. Rehmane, and D. Culler, "A living laboratory study in personalized automated lighting controls," in *Proceedings of the third ACM workshop on embedded sensing systems for energy-efficiency in buildings*. ACM, 2011, pp. 1–6.
- [18] X. Jiang, M. Van Ly, J. Taneja, P. Dutta, and D. Culler, "Experiences with a high-fidelity wireless building energy auditing network," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2009, pp. 113–126.
- [19] B. Balaji, C. Verma, B. Narayanaswamy, and Y. Agarwal, "Zodiac: Organizing Large Deployment of Sensors to Create Reusable Applications for Buildings," in *Proceedings of the ACM Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM, 2015.
- [20] A. Bhattacharya, D. Culler, D. Hong, K. Whitehouse, J. Ortiz, and E. Wu, "Automated Metadata Construction To Support Portable Building Applications," in *Proceedings of the ACM Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM, 2015.
- [21] J. Gao, J. Ploennigs, and M. Berges, "A Data-driven Meta-data Inference Framework for Building Automation Systems," in *Proceedings of the ACM Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM, 2015.
- [22] D. Hong, H. Wang, J. Ortiz, and K. Whitehouse, "The Building Adapter: Towards Quickly Applying Building Analytics at Scale," in *Proceedings of the ACM Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM, 2015.
- [23] M. Pritoni, A. Bhattacharya, D. Culler, and M. Modera, "Short Paper: A Method for Discovering Functional Relationships Between Air Handling Units and Variable-Air-Volume Boxes From Sensor Data," in *Proceedings of the ACM Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM, 2015.
- [24] J. Weimer, S. A. Ahmadi, J. Araujo, F. M. Mele, D. Papale, I. Shames, H. Sandberg, and K. H. Johansson, "Active actuator fault detection and diagnostics in hvac systems," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 107–114.
- [25] M. Padilla and D. Choinière, "A combined passive-active sensor fault detection and isolation approach for air handling units," *Energy and Buildings*, vol. 99, pp. 214–219, 2015.

- [26] N. Fernandez, M. R. Brambley, and S. Katipamula, *Self-correcting HVAC Controls: Algorithms for Sensors and Dampers in Air-handling Units*. Pacific Northwest National Laboratory, 2009.
- [27] M. Padilla, D. Choinière, and J. A. Candanedo, "A model-based strategy for self-correction of sensor faults in vav air handling units," *Science and Technology for the Built Environment*, no. just-accepted, pp. 00–00, 2015.
- [28] D. Hong, J. Ortiz, K. Whitehouse, and D. Culler, "Towards automatic spatial verification of sensor placement in buildings," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 2013, pp. 1–8.
- [29] B. Balaji, H. Teraoka, R. Gupta, and Y. Agarwal, "Zonepac: Zonal power estimation and control via hvac metering and occupant feedback," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 2013, pp. 1–8.
- [30] M. Hydeman, *Advanced Variable Air Volume: System Design Guide: Design Guidelines*. California Energy Commission, 2003.
- [31] "VAV Terminal Control Applications Application Note," *Johnson Controls Technical Bulletin*, 2003.
- [32] S. T. Bushby, "BACnet™: A standard communication infrastructure for intelligent buildings," *Automation in Construction*, vol. 6, no. 5, pp. 529–540, 1997.
- [33] T. Weng, A. Nwokafor, and Y. Agarwal, "Buildingdepot 2.0: An integrated management system for building analysis and control," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 2013, pp. 1–8.
- [34] "scikit-learn Machine Learning in Python," 2015, <http://scikit-learn.org/>.
- [35] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *KDD workshop*, vol. 10, no. 16. Seattle, WA, 1994, pp. 359–370.
- [36] R. S. Tsay, *Analysis of financial time series*. John Wiley & Sons, 2005, vol. 543.
- [37] A. Foucquier, S. Robert, F. Suard, L. Stéphan, and A. Jay, "State of the art in building modelling and energy performances prediction: A review," *Renewable and Sustainable Energy Reviews*, vol. 23, pp. 272–288, 2013.
- [38] F. V. Jensen, *An introduction to Bayesian networks*. UCL press London, 1996, vol. 210.
- [39] J. Pearl, *Causality*. Cambridge university press, 2009.
- [40] L. Ljung, *System identification*. Springer, 1998.
- [41] K. Godfrey, *Perturbation signals for system identification*. Prentice Hall International (UK) Ltd., 1993.
- [42] K. Godfrey, H. Barker, and A. Tucker, "Comparison of perturbation signals for linear system identification in the frequency domain," in *Control Theory and Applications, IEE Proceedings-*, vol. 146, no. 6. IET, 1999, pp. 535–548.
- [43] K. Godfrey, A. Tan, H. Barker, and B. Chong, "A survey of readily accessible perturbation signals for system identification in the frequency domain," *Control Engineering Practice*, vol. 13, no. 11, pp. 1391–1402, 2005.
- [44] J. Teeter and M.-Y. Chow, "Application of functional link neural network to hvac thermal dynamic system identification," *Industrial Electronics, IEEE Transactions on*, vol. 45, no. 1, pp. 170–176, 1998.
- [45] J. Sousa, R. Babuška, and H. Verbruggen, "Fuzzy predictive control applied to an air-conditioning system," *Control Engineering Practice*, vol. 5, no. 10, pp. 1395–1406, 1997.
- [46] G. Virk, J. Cheung, and D. Loveday, "Practical stochastic multivariable identification for buildings," *Applied mathematical modelling*, vol. 19, no. 10, pp. 621–636, 1995.
- [47] X.-F. Liu and A. Dexter, "Fault-tolerant supervisory control of vav air-conditioning systems," *Energy and Buildings*, vol. 33, no. 4, pp. 379–389, 2001.
- [48] S. Wang and Y. Chen, "Fault-tolerant control for outdoor ventilation air flow rate in buildings based on neural network," *Building and Environment*, vol. 37, no. 7, pp. 691–704, 2002.
- [49] R. Fontugne, J. Ortiz, D. Culler, and H. Esaki, "Empirical mode decomposition for intrinsic-relationship extraction in large sensor deployments," in *Workshop on Internet of Things Applications, IoT-App*, vol. 12, 2012.
- [50] S. Dawson-Haggerty, A. Krioukov, J. Taneja, S. Karandikar, G. Fierro, N. Kitaev, and D. E. Culler, "Boss: Building operating system services," in *NSDI*, vol. 13, 2013, pp. 443–458.
- [51] "NiagaraAX by Tridium," 2015, <http://www.niagaraax.com/>.